

REMARKS/ARGUMENTS

Claims

Claims 1-6, 8-14, 16, 18, 20-26 and 31-38 are currently pending in this application. Claims 1, 2, 14, 18, 32, 35, and 37 are currently amended. Independent Claim 1 has been amended only to improve readability. Claim 3 has been amended to clarify that "some" means at least two. Claims 14, 18, 32, 35, and 37 have been amended to correct informalities.

CLAIM OBJECTIONS

Claims 18, 32, 35, and 37 have been objected to because of informalities. The Examiner suggested replacing "processing nodes" with "processor nodes" to make the claim language consistent. (OA, p.2.) The Applicants have replaced "processing nodes" with "processor nodes" in Claims 18, 32, 35, and 37 to as suggested by the Examiner and believes the objections to these claims have been overcome.

The Examiner has objected to Claim 14 "as being of improper dependent form for failing to further limit the subject matter of a previous claim." (OA, p. 3.) Applicants have added the limitation that the interface is "configured to receive the additional previously presented instructions." The Applicants believe this additional limitation overcomes the Examiner's objection.

REJECTIONS UNDER 35 U.S.C. §103

Claims 1-6, 8-14, 16, 18, 20-26 and 31-38 stand rejected under 35 U.S.C. §103(a) as being unpatentable over Borkar et al., "iWarp: An Integrated Solution to High-Speed Parallel Computing" (Hereinafter Borkar) in view of Barat et al.,

"Reconfigurable Instruction Set Processors: A Survey" (Hereinafter Barat). The Applicants respectfully traverse these rejections.

The Examiner has correctly stated that "Borkar has not taught that the plurality of processor nodes additionally comprise a software extensible device." (Office Action p.4.) However, the Examiner has asserted that "Barat has taught coupling a reconfigurable processing unit (RPU) to a microprocessor," and that it would have been obvious "to modify the iWarp cells of Borkar to each comprise a software extensible device (i.e. an RPU) coupled to the computation unit of each iWarp cell."

The Examiner has further asserted that a "suggestion/motivation for [combining Barat and Borkar] would have been that processing is more specialized thereby accelerating execution and increasing performance [section 1, "introduction"]." (Office Action p. 5.) The Applicants respectfully submit that the "suggestion/motivation" asserted by the Examiner is neither a suggestion nor a motivation, there is no expectation of success, and that Barat teaches away from the combination.

First, there is no motivation to combine the RPU of Barat with the iWarp of Borkar. Borkar achieves acceleration in speed by increasing the number of processors in the array. "a large array of iWarp cells will deliver an enormous computing bandwidth never before realized in distributed memory parallel systems." (Borkar, section 1., third paragraph.) "iWarp is intended for systems of various sizes ranging from several processors to thousands of processors." (Borkar, section 1., fifth bullet.)

Further, Borkar achieves "specialized processing" using special-purpose arrays and achieves high speed in the special purpose arrays using large numbers of iWarp cells. "With the iWarp cell, various special-purpose arrays that execute only a predetermined set of these algorithms can easily be built. . . .

In areas such as high-speed signal processing, special purpose arrays can effectively use hundreds or even thousands of iWarp cells.” (Borkar, section 3.2) Thus, there is no need to add an RPU to the iWarp for accelerating execution, specializing processing, or increasing performance.

Second, even if the RPU of Barat were combined with Borkar, the Examiner has not established an expectation of success in “accelerating execution and increasing performance.” Barat states that there is not always an improvement from adding an RPU to a processor. For example, Barat states that:

“The benefit obtained from executing a piece of code in the RPU depends on communication and execution costs [10]. The time needed to execute an operation in the RPU is the sum of the time needed to transfer the processed data and the time required to process it. If this total time is smaller than the time it would normally take in the processor alone, then an improvement can be obtained.” (Barat, Section 2.1, first paragraph.)

The Examiner has not set forth a comparison of the “time needed to execute an operation in the RPU” with “the time it would normally take in the [iWarp] alone,” or any other reasoning that would indicate an expectation of a benefit of using the RPU of Barat in an array of iWarp devices of Borkar. The improvement in performance by increasing the number of processors in the array and/or creating specialized arrays in Borkar does not depend on communication and execution costs.

Third, Barat teaches away from the combination. The iWarp uses long instruction words to achieve a high computation rate. However, Barat states that “[RPU] designs for VLIW processors would normally involve fixed duration instructions. Unknown duration on VLIW would result in pipeline stalls; with a major loss in performance. *No VLIW processors currently exist that include reconfigurable hardware.*” (Barat, section 2.5 third paragraph.) (Emphasis added)

Whereas Borkar states:

... iWarp can attain a high computation rate consistently ...
because the multiple functional units in the computation agent are
directly accessible through a long instruction word (LIW)
instruction. By translating user's code directly into these long
instructions using an optimizing compiler [17], a high computation
rate can be achieved for all programs, vectorizable or not.
(Borkar, section 6.)

Thus, according to Barat, combining the RPU of Barat with the iWarp of
Borkar would result in pipeline stalls, with a major loss in performance.

Finally, the Barat reference is a survey article titled "*Reconfigurable
Instruction Set Processors: A Survey*" which Barat asserts is thorough. However,
there is no teaching or suggestion that in the survey article of using an RPU in a
processor array or with multiple processors. Moreover, while Barat makes
suggestions for future work, none of the suggestions include combining an RPU
with an array of processors. In the conclusion, Barat states:

The two main aspects that have to be studied are the interfacing of
the reconfigurable logic with the microprocessor and the design of
reconfigurable logic itself. They both involve many decisions,
which we have tried to enumerate as thoroughly as possible.

Future work will continue along these main lines. Both the
hardware and the software side have to be studied further. Several
experiments will have to be done in order to determine which is the
best RISP architecture for the MPEG-4 domain. Similarly, during
the selection of the architecture, tests will have to be done to
determine the best way to compile code.
(Barat, section 4.)

For at least the above reasons, Examiner has not met the first element of a
prima facie case for obviousness as there is no suggestion or motivation, either in
the references themselves or in the knowledge generally available to one of
ordinary skill in the art, to modify or combine references.

Even if Borkar and Barat were combined, they do not teach all the elements of Claim 1. Borkar teaches a processor component called an "iWarp." The iWarp cell includes an iWarp component and a local memory. Multiple iWarp cells are connected in an array of cells or nodes. For example, nine iWarp cells may be connected to each other in a three by three torus array. (see Figure 5 of Borkar.)

While the terms are similar, the "input ports" and "output ports" have a different meaning and different use in Borkar from the term "standard input/output interface" of Claim 1 of the present invention. The term "input port" or "output port" in Borkar means an interface for inter-cell communication, whereas a similar term "standard input/output interface" in Claim 1 means an interface for communication with a peripheral device. As correctly pointed out by the Examiner, "each iWarp cell has 4 input ports and 4 output ports configured to interface to *other iWarp cells*." (Office Action, p. 4.) (Emphasis added) However, all four input and output ports are for inter-cell communication with neighboring iWarp cells, not standard input/output communications.

Borkar uses the local memory device on the iWarp cell for communication with peripherals external to the array instead of using an input port or an output port. "[T]he iWarp cell's connection with peripherals uses the local memory, while its intercell connection uses ports of iWarp." (Borkar, Section 2.2, third paragraph.) However, a memory used for memory mapped I/O is not the same as a "standard input/output interface" in Claim 1.

In rejecting Claim 1, the Examiner asserts that Borkar has taught "a **first** standard input/output interface [*first output port*] configured to communicate with a first input/output device [*The communication agent of each iWarp cell has 4 input ports and 4 output ports configured to interface to other iWarp cells; section 2.1*]"

(Bold emphasis added.) (Italics in original.) The Examiner further asserts that Borkar has taught "a **second** standard input/output interface [*first output port*] configured to communicate with a second input/output device [section 2.1]." (Office Action, p. 4.) (Bold emphasis added.) (Italics in original.) The Applicants respectfully submit that the Examiner is in error.

First, as discussed above, the terms "input port" and "output port" as used in Borkar refer to interfaces used for cell to cell communications, not peripheral communication. Thus, the interface that the Examiner refers to in Borkar as "*4 input ports and the 4 output ports*" is used for cell to cell communications, not communication with external peripheral devices:

"There are two ways that an iWarp cell, consisting of an iWarp component and its local memory, physically interfaces with the external world. Recall that the iWarp component has four input ports and four output ports. The first interface method is to use a physical bus to connect an output port of an iWarp to an input port of another [iWarp]. . . . Usually another unidirectional bus in [an] opposite direction is also provided, so that bidirectional data communication between the two [iWarp] component[s] is possible." (Borkar, Section 2.2, second paragraph.)

Therefore, the [*first output port*] referred to in Borkar is for cell to cell communication and is neither the "first standard input/output interface" nor the "second standard input output interface," that are recited in Claim 1.

Second, Borkar uses only one memory device in each iWarp cell for memory mapped I/O. Borkar uses the local memory for peripheral communication (second interface method) instead of the input port and output ports (first interface method). In Borkar:

"The second interface method is via the local memory of the iWarp cell, as depicted by Figure 4.

Using this interface the iWarp cell can reach peripherals such as standard busses, disks, graphics devices and sensors.” (Borkar, Section 2.2, third paragraph.)

However, only one local memory appears in the iWarp cell as depicted in Figure 4 and Borkar discloses only one local memory in each iWarp. Thus, no iWarp cell includes a second interface for communication with peripheral devices in Borkar. (See Figures 3 and 4.) Therefore, there is no second interface in the iWarp of Borkar. Moreover, there is no need for a second interface since the local memory in each individual cell can be used as an interface. “[P]eripherals can be attached to any set of iWarp cells in an array of iWarp cells.”

Third, the iWarp memory of Borkar is not a “standard input/output interface configured to communicate with [an] input/output device” that is recited in Claim 1. Memory mapped I/O requires an interface to translate the address and data written to a memory address into a standard I/O protocol. Memory devices do not perform standard I/O with peripheral devices. While the memory device is on the iWarp, there is no teaching in Borkar of a standard input/output interface on the iWarp.

Since neither Borkar nor Barat teach “a first standard input/output interface configured to communicate with a first input/output device . . . and a second standard input/output interface configured to communicate with a second input/output device,” as recited in Claim 1, the combination of Borkar and Barat likewise does not teach these elements of Claim 1.

For at least the above reasons, the Applicants believe that Claim 1 is allowable. Applicants also believe that Claims 2-6, 8-14, 16, and 31-36 are allowable for at least the reasons discussed with respect to Claim 1, from which they depend.

Claim 18

Claim 18 recites in part "determining if a neighboring device is a member of the plurality of processor nodes."

The Examiner has asserted *inter alia* that "[s]ince communication depends upon the neighboring device, there is inherently a determination if the neighboring device is a member of the plurality of iWarp cells; section 2.2."

(Office Action, p. 8.) The Applicants respectfully submit that the Examiner is in error.

First, the fact that an operation *may* be performed does not establish inherency of the operation. "[t]he fact that a certain result or characteristic may occur or be present in the prior art is not sufficient to establish the inherency of that result or characteristic." In re Rijckaert, 9 F.3d 1531, 1534. Moreover, "[t]o establish inherency, the extrinsic evidence must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill." In re Robertson, 169 F.3d 743, 745, 49

Therefore, to establish inherency of determining "if a neighboring device is a member of the plurality of processor nodes," the Examiner must show that the determination must necessarily be performed for the communication with the neighboring device, not that such determination may happen or even may be useful. If a node can communicate with a neighboring node without the determination then it is not inherent. The Examiner has not shown that the determination must necessarily be performed. Moreover, the iWarp of Borkar can and does communicate with neighboring iWarp cells without determining if the neighbor is an iWarp cell. For example, in spooling communication, "the sender process can simply specify the destination and message location and continue with its processing." (Borkar, Section 4.1.1 "Spooling")

Second, in communication between cells, the iWarp does not depend upon the identity of the neighboring device. In Borkar, determining "if a neighboring device is a member of the plurality of processor nodes" is not necessary for communication because all communications using the ports of the iWarp are to another processor node. Moreover, as discussed above, ports are used only for inter-node communications. Borkar teaches that each iWarp is hardwired via ports to neighboring iWarp cells. An iWarp may select a neighboring iWarp cell for communication by selecting a port. Moreover, any port selected will necessarily result in a communication with another iWarp cell. Selecting a neighboring iWarp is not the same as determining if the selected cell is an iWarp cell. Therefore, the iWarp can send a signal to another node by selecting a port without determining that the port is connected to an iWarp cell.

The iWarp ports of Borkar are connected only to neighboring iWarp nodes. The peripherals are connected to the iWarp only via local memory. The connections between nodes are independent of the connections to peripherals. "[t]he iWarp cell's connection with peripherals uses the local memory, while its intercell connection uses ports of iWarp. Since these two functions use different physical resources of the iWarp cell, they can be implemented in-dependently from each other." (Borkar, Section 2.2, last paragraph.) Therefore, any communication sent or received via a port is to a neighboring iWarp node. Moreover, any communication a peripheral will use the local memory of the iWarp not a node. See, for example, Figure 5 illustrating a 3X3 torus array in which each cell, (node) has four pairs of port connections and each of the four pairs of port connections are connected to another cell. Regarding Figure 5 Borkar states that, "[p]eripherals can be attached to any of the iWarp cells via its local memory." (Borkar, Section 3.1, first paragraph.) Thus, it is not necessary to

determine if the neighboring device is a member of the plurality of nodes in order to direct communication to either a node or a peripheral device.

For at least the reason that the missing element of "determining if a neighboring device is a member of a plurality of processor nodes" is not inherent, the Applicants believe that Claim 18 and those claims that depend therefrom are allowable.

The Examiner has correctly stated that "Borkar has not taught providing an additional new instruction to a set of standard instructions for the processing element, using at least one software extensible device in the plurality of the processor nodes, wherein the new instructions can be programmed by software." (Office Action p. 9.) The Examiner has repeated the arguments with respect to Claim 1 for combining Barat with Borkar. For at least the reasons discussed above with respect to Claim 1 the Applicants believe that the Examiner has not provided a prima facie case of obviousness for Claim 18.

The Examiner suggests no teaching in Barat of a "determining if a neighboring device is a member of the plurality of processor nodes," as recited in Claim 18 and the Applicants find no such teaching in Barat. Therefore the combination of Borkar and Barat does not teach all the elements of Claim 18.

For at least the above reasons, the Applicants believe that Claim 18 is allowable. Applicants believe that Claims 20-26 and 37-38 are allowable for at least the reasons discussed with respect to Claim 18 from which they depend.

Claims 10 and 22

The Examiner has asserted that "[r]eferring to claims 10 and 22, taking claim 10 as exemplary, Keller and Barat have taught . . ." (Office Action p.6.) (emphasis added.) The Applicants are unable to find a citation to the reference of "Keller" in the Office Action or any of the lists of references cited by the

Examiner in previous Office Actions. The Applicants respectfully suggest that "Keller" may be an incorrect reference citation due to a typographical error and requests the Examiner to identify the correct reference.

Claims 11 and 23

The Examiner has asserted that "[r]eferring to claims 11 and 23, taking claim 11 as exemplary, Keller and Barat have taught . . ." (Office Action p.6.) (emphasis added.) The Applicants are unable to find a citation to the reference of "Keller" in the Office Action or any of the lists of references cited by the Examiner in previous Office Actions. The Applicants respectfully suggest that "Keller" may be an incorrect reference citation due to a typographical error and requests the Examiner to identify the correct reference.

Applicants believe that all pending claims are allowable and respectfully requests that the Examiner issue a Notice of Allowance. Should the Examiner have questions, Applicants' undersigned representative may be reached at the number provided below.

Respectfully submitted,

Ricardo E. Gonzalez et. al.

Date: November 5, 2007



Ronald Rohde, Reg. No. 45,050

Carr & Ferrell *LLP*

2200 Geng Rd.

Palo Alto, CA 94303

Phone (650) 812-3487

Fax (650) 812-3444